

**Abstract**

1        A technique for efficiently populating a cache in a data processing system with resources is  
2 disclosed. In particular, a node in accordance with the illustrative embodiment of the present  
3 invention defers populating its cache with a resource until at least two requests for the resource have  
4 been received. This is advantageous because it prevents the cache from being populated with  
5 infrequently requested resources. Furthermore, the illustrative embodiment of the present invention  
6 populates a cache with a resource only when: at least  $i$  requests for the resource have been received at  
7 a given node within an elapsed time interval,  $\Delta t$ , wherein  $i$  is an integer greater than one; and at least  
8 one request for the resource has been received from at least  $n$  of the  $m$  filial nodes of the given node  
9 within an elapsed time interval,  $\Delta t$ , wherein  $m$  is an integer greater than one,  $n$  is an integer greater  
10 than one, and  $m \geq n$ . Embodiments of the present invention are particularly advantageous in computer  
11 networks that comprise a *logical* hierarchical topology, but are useful in any computer network, and in  
12 individual data processing systems and routers that comprise a cache memory.